A BRIEFING FOR CX AND CONTACT CENTER LEADERSHIP

No Agent Left on an Island

How to manage, coach and support the agents who represent your brand, including the thousands of them you do not employ and have never seen work.

Written for the brands that buy contact center services. Not for the BPOs that deliver them.

Cloud Tech Gurus · CX Consulting and Advisory Firm

August 2026 · Vendor-neutral, no supplier named in this document

Prepared for organizations operating 200 to 10,000+ agents, internal or outsourced

EXECUTIVE SUMMARY

The floor did more than you think it did.

A contact center floor was never just a room. It was a support system. A new hire could turn around and ask a question. A supervisor noticed the agent who went quiet at two in the afternoon. A team lead saw the screen, saw the face, and helped before a bad call became a bad week. None of that was in a system. It was in the proximity.

Most of that proximity is gone. Agents work from homes and satellite sites, and a large share of them work for a company you do not own, in a country you have never visited. The work still gets done. What went missing is the layer that kept the people doing it supported, coached and connected.

Why it matters

Your agents represent your brand to your customers whether they are on your payroll or your partner's. An agent who cannot get an answer for forty minutes gives a worse answer. An agent who is struggling and unnoticed becomes an escalation, then an attrition statistic. And a supervisor who cannot see the work cannot coach it, which means your quality program is reduced to describing problems after the customer has already experienced them.

For outsourced seats there is a second issue layered on top. Almost everything you know about that floor arrives in a report written by the party you are evaluating. That is not an accusation. Well run partners report honestly. It is a structural limit: a summary written after the fact cannot substitute for seeing the work, and no clause in an agreement makes an unobservable rule enforceable.

> **The question is not how to watch your partner more closely. It is how to stand next to agents you cannot stand next to.**

> **How this document is written**
>
> Cloud Tech Gurus is compensated by technology suppliers, not by the organizations we advise. That model only survives if our recommendations hold up under scrutiny. So this briefing names no vendor, quotes no marketing language, and states the boundaries of its own evidence in Section 7. The evaluation rubric in Section 8 is written so that you can use it against us.

Contents

1	What the floor actually did	4
2	The outsourced version of the same problem	6
3	Where your program stands: a four tier model	8
4	Why more reporting will not fix it	10
5	What a virtual floor actually is	12
6	The compliance layer, and where liability sits	14
7	What the evidence shows, and its limits	16
8	How to evaluate this category: ten questions	18
9	What we recommend, and why we are saying it	20

Who this is for

Enterprises and brands that run customer contact operations, whether those agents sit on your payroll, on a partner's, or both. Specifically: Vice Presidents and Directors of contact center operations, Chief Customer Officers, heads of customer care, vendor and outsourcing managers, and the risk, compliance, finance and procurement stakeholders who approve their decisions.

Who this is not for

This is not written for BPOs or outsourcing providers. It is written for their clients. If you deliver contact center services rather than buy them, the perspective here is the one your clients hold, which may still be worth reading, but nothing in this document is addressed to you.

What it is not

It is not a product brochure, and it is not a case against outsourcing. Outsourcing works. This briefing is about the support and coaching layer that thinned out when floors went distributed, and what can now be done about it.

> **The two minute version**
>
> Ask your BPO to let you spend one unscheduled hour with six agents on your program. Faces, screens, live call audio. Not to catch anyone out. To see the work.
>
> Most partners cannot accommodate that, and not because they refuse. The capability does not exist on their floor. Everything in this document follows from that gap.

SECTION 1

What the floor actually did

Before discussing technology it is worth being precise about what was lost, because the answer is not "visibility." Visibility is a means. Here is what the physical floor delivered, none of which anyone wrote down as a requirement because nobody had to.

Answers in seconds

An agent stuck mid call turned their head. A tenured colleague two seats over knew the answer. The customer waited fifteen seconds instead of being placed on hold, transferred, or promised a callback that generated a second contact.

Noticing

Supervisors did not audit their teams. They walked past them. They saw the agent who had stopped smiling, the one who was fighting the same system error repeatedly, the one whose posture had changed since the morning huddle. Almost all early intervention in a contact center came from someone noticing something.

Coaching in the moment

The most effective coaching a contact center ever produced happened during the interaction or immediately after it, while the agent still remembered the call and cared about it. A scorecard delivered three weeks later is a record, not coaching.

Learning by proximity

New hires sat beside experienced agents and absorbed the job by listening. Nesting worked because it was physical. There was a gradient between training and going solo, and the gradient was made of the person sitting next to you.

Belonging

Contact center work is hard, repetitive and emotionally demanding. What made it survivable for many people was the team around them. Culture was not a program. It was a room.

Section 1: What the floor actually did

What distance costs

Remove proximity and every one of those five functions degrades, usually quietly. The metrics move before anyone identifies why.

On the floor	Distributed, without a substitute
Turn and ask	Post in a channel and wait. Average resolution time measured in tens of minutes, often resolved by the customer hanging up.
Supervisor notices	Supervisor learns from a metric, days later, after the pattern is established.
Coaching during the call	Coaching in a scheduled one to one, referencing a call the agent has forgotten.
Nesting beside a peer	Shadowing in a conference call where only one person can share a screen. Then solo.
Team around you	A queue, a headset, and a kitchen table.

The compensating behaviour is predictable and counterproductive. Supervisors add check ins, stand ups, status pings and reporting because it is the only lever they have. Agents experience that as micromanagement. Supervisors experience it as flying blind. Both readings are correct, and the underlying cause is the same: the mechanism that used to do this work was physical, and nothing replaced it.

> **Nobody decided to stop supporting agents. The method for supporting them was proximity, and proximity is what remote work removed.**

Why this is not a nostalgia argument

Very few operators want to go back. Distributed work expanded hiring pools, improved retention for many roles, and cut real estate cost. The argument here is not that the floor should return. It is that the floor was doing five jobs, and most organizations replaced only the two that were easy to see: the phone system and the reporting.

SECTION 2

The outsourced version of the same problem

Everything in Section 1 applies with more force to outsourced seats, because there the distance is organizational as well as physical.

Consider a brand running 180 internal agents and 1,400 outsourced seats across three BPO partners in two countries. The support and coaching asymmetry between those two populations is usually stark, and it is rarely deliberate.

Function	Internal agents	Outsourced agents
Who supports them day to day	Your supervisors	The vendor's supervisors
Who coaches to your brand standard	You do	Someone you have never met
Real time help	Same shift, often same call	Escalation path measured in days
Quality sampling	You choose the sample	The vendor selects and submits it
Coaching visibility	You see it happen	You see it summarized in a review
Policy adherence	Observed	Attested to in writing
Attrition insight	Leading indicators	Reported after the fact, often netted
Problem discovery	Minutes to hours	Customer escalation, or next review

The last row is the expensive one. When something goes wrong on an outsourced floor, the reporting cycle is rarely how you find out. You find out from a customer complaint that reaches an executive, from a regulator, from a social post, or from unexplained attrition on your program.

Every row above describes something that exists internally and has no outsourced equivalent. Not because anyone chose to drop it, but because the mechanism that delivered it was proximity, and outsourcing trades proximity away by design.

Section 2: The outsourced version of the same problem

The three input problem

Nearly all outsourced insight rests on three inputs: the quarterly review deck, the performance scorecard, and the call samples. Each is genuinely useful. All three share one origin.

In audit terms this is a population with a single interested source and no independent verification path. In an internal control review it would be flagged immediately. In outsourcing it is simply how the industry has always worked, largely because visiting the floor was the only alternative, and that meant an international flight, scheduled in advance, on a day the vendor knew you were coming.

The cost of a scheduled look

Site visits have real value and we are not arguing against them. But by construction they are the least representative sample available: announced, prepared for, and conducted on the vendor's best day. Two managers on a plane twice a year is a governance ritual, not a support model.

Say what you actually want

There is a version of this conversation that goes badly with partners, and it starts with the word audit. There is another version that goes well, and it sounds like this: those agents represent our brand to our customers, we would like to help coach and support them the way we coach and support our own, and we would like our supervisors and yours working from the same picture.

That is a different request, and in our experience it gets a different answer. It is also closer to what most buyers actually want. Very few operators are trying to catch their BPO out. Most are trying to make a population of agents better, and have no mechanism for doing it.

> **Before your next review**
> Ask for one unscheduled hour with six agents. Note the answer and the reason given. If the reason is that the tooling does not exist, you have found the gap precisely, and it is not a partner problem. It is a capability problem, and it is solvable.

SECTION 3

Where your program stands

Four tiers of distributed and outsourced support maturity. Most enterprise programs we assess sit in Tier 1 or Tier 2, including sophisticated operators with mature internal CX functions. Score yourself honestly. Your tier is set by what you can actually do, not by what your agreement entitles you to.

Tier 1: Reported
You receive a scorecard and a review deck. Call samples are vendor selected. You have no view of daily work and no ability to help. Problems surface through customer escalation. *Support capability: none.*

Tier 2: Sampled
You pull your own recordings, run your own quality reviews on a sample, and may have dashboard access to the vendor's reporting. Better, but still retrospective and still a sample, typically a fraction of one percent of interactions reviewed days later. *Support capability: after the fact.*

Tier 3: Present
You can see the floor as it works: agent presence, screen activity and live call audio, on demand, without scheduling it. You can help an agent during the interaction. Problems surface as they occur rather than in a report. *Support capability: real time.*

Tier 4: Integrated
Outsourced and internal agents work in the same rooms. Your supervisors and the vendor's coach the same population against the same standard. Training and nesting are shared. The distinction between internal and outsourced becomes contractual rather than experiential. *Support capability: continuous, and shared.*

A program can sit in different tiers with different partners, and usually does. Score your largest partner first, then your newest.

The jump that actually matters

Tier 1 to Tier 2 improves the quality of your evidence. You learn more about what already happened. Useful, and it will make your reviews more honest.

Tier 2 to Tier 3 is a change in kind. It moves you from describing the past to affecting the present. Every performance gain documented in Section 7 comes from that transition. Not from better reporting, but from a supervisor being able to help while helping still changes the outcome.

Why Tier 4 is the one to aim at

The operators getting the most out of this are not running an audit function. They put outsourced and internal agents in the same rooms, let both supervisor populations coach against one standard, and run shared nesting for new hires on both sides. Better oversight is a byproduct. The performance benefit comes from the outsourced team being treated as part of the floor rather than as a supplier deliverable.

That reframing also resolves most of the partner relations objection, which is worth stating plainly because it is the objection that stalls these projects. A brand asking to audit its BPO creates friction. A brand asking to help coach, train and support its BPO's agents alongside its own is offering something the BPO wants, at no cost to the BPO, and generally gets a different answer.

<1%

Share of interactions reviewed under a typical quality sampling regime

90 days

Common review to review interval during which issues compound unnoticed

1

Independent sources verifying your outsourced performance data

SECTION 4

Why more reporting will not fix it

The instinct when insight feels thin is to ask the partner for more: additional metrics, a tighter scorecard, an extra review cadence, a third party auditor, sometimes a second BPO positioned to keep the first honest. Each has merit. None closes the gap, for five structural reasons.

1. The source does not change

A richer report from the same party is still a self assessment. Adding fields to a scorecard increases resolution, not independence. If the underlying observation is made by the party being measured, the measurement inherits that party's interests no matter how granular it becomes.

2. Sampling cannot catch what it does not sample

Quality programs typically review well under one percent of interactions. That is adequate for detecting broad trends and inadequate for detecting a single agent working from an unauthorized location, a policy breach on one shift, or a coaching failure confined to one team lead. The problems that generate regulatory and reputational exposure are usually specific, not distributional.

3. Retrospective data cannot be acted on

A finding delivered thirty days after the interaction has no remaining operational value. The customer has already had the experience, the agent has repeated the behaviour several hundred times, and the coachable moment is gone. Latency is not an inconvenience in this context. Past a certain point it makes the data ornamental.

The two responses that look like solutions

4. Adding a third party adds a layer, not a line of sight

An external auditor, or a second BPO engaged to review the first, is still working periodically, from the incumbent's systems and the incumbent's records, on a scheduled visit. You have purchased a second opinion about the same evidence, plus a new commercial relationship with its own incentives.

In our experience this reliably increases cost and rarely increases certainty. It is also the arrangement most likely to damage the partnership, because it signals distrust without giving either side better information. If the incumbent learns that a competitor is reviewing them, you have created a political problem on top of an unsolved operational one.

5. Unobservable obligations are unenforceable obligations

Your agreement almost certainly prohibits mobile phones at the workstation, requires a clean desk where sensitive data is handled, and bars any unauthorized person from seeing the screen or hearing the call. On a physical floor a supervisor walking past enforced all three. In a distributed model they are enforced by attestation.

An attestation is a promise, not a control. A promise is also not something you can produce for a regulator, which is the subject of Section 6.

The distinction that matters

None of this argues for distrusting your partner. It argues that **trust is not a control, and it is not a support model either**. Mature outsourcing relationships are strengthened, not strained, when both parties work from the same live picture instead of negotiating over whose report is right. In our experience the operators who adopt this most readily are the ones whose partner relationships are already healthy.

SECTION 5

What a virtual floor actually is

This category is frequently mislabelled as monitoring software, including by people who sell it. That label describes a different product and predicts a different outcome. What the credible platforms in this category are attempting is narrower and more useful: to reconstruct the five functions in Section 1 for people who are not in the same room.

Rooms, not sessions

Agents work together in a persistent production room, the way a team occupies a section of floor. They can see each other. They can see which supervisors are present. Team rooms, huddle rooms and training rooms exist alongside it, and moving an agent between them takes a few clicks rather than a calendar invitation.

Faces and screens, together, for everyone

A supervisor sees each agent's video and up to three of their monitors at the same time, across the whole team, in one view. This is the technical distinction that makes the category possible. General purpose conferencing tools show many faces but permit one screen share, which is why using them for this has never worked.

Help that goes both directions

The agent can raise a hand, message a supervisor, or call a colleague inside the room and get an answer in seconds, without leaving the call. This is the part most often overlooked in evaluations and the part that determines whether agents accept the tool or resent it. A one directional system is a surveillance system, and it will fail on adoption.

Coaching during the work

A supervisor can listen to the live customer conversation while watching the agent's screens, and guide them then and there. Not "you handled that badly three weeks ago," but "you are in the wrong system, try the other tab."

Nesting, noticing, and what is deliberately absent

Nesting that resembles nesting

A new hire can shadow a tenured agent: listening to the live call while watching that agent's monitors. This restores the gradient between classroom and going solo that distributed hiring removed, and it is the feature operators tell us they underestimated.

Noticing, at scale

These platforms can surface when an agent appears to be struggling, based on facial expression state, so a supervisor can check in. Read that carefully, because the framing matters and the wrong framing will get the project rejected: the purpose is to replace the supervisor who used to glance across the room and see that someone was having a rough morning. Used as a performance metric it is indefensible. Used as a prompt to go help a person, it is the closest software has come to reproducing what a good team lead did naturally.

What is deliberately absent

The credible platforms in this category **do not record**. No video is stored, no screen capture is retained, no personal data is collected. Alerts go to a supervisor in the moment and nothing persists afterward.

That is not a limitation, it is the design. It makes deployment possible under employee privacy regimes and works council review that would reject a recording system outright, and it is the line between a management tool and surveillance. A platform that records creates a discovery liability, a retention policy question and a workforce relations problem on the day it goes live. Get the answer in writing.

> **On the word monitor**
>
> One supplier in this category refuses the word entirely. Their position, quoted with permission and without attribution: monitoring is a terrible description of what this is, because the goal is to replicate a floor, and floors are for managing, supporting and coaching people.
>
> We repeat it here because vocabulary predicts outcomes. Programs introduced to agents as monitoring fail. Programs introduced as putting the team back in the same room, with the supervisor available instead of absent, succeed. The technology is identical. The framing is not.

SECTION 6

The compliance layer

For organizations handling card data, health records or other regulated information through distributed or outsourced seats, there is a second argument that runs parallel to everything above. It is a liability argument, and the liability is asymmetric.

Where the exposure actually sits

When a customer's data is compromised by an agent working on your program, the notification obligation is generally yours. The letter goes out under your brand. The regulatory correspondence names your organization. The press coverage uses your logo, not your vendor's.

Meanwhile most outsourcing agreements cap vendor liability at a defined multiple of fees, frequently a small fraction of the total cost of a notification event once legal, forensic, remediation, credit monitoring and churn are accounted for. We would encourage you to read your own indemnity and limitation of liability clauses with that comparison in mind. This briefing is not legal advice, and the exercise takes ten minutes.

Three rules that became unverifiable

Policy	How it was enforced on a floor	How it is enforced now
No mobile devices at the workstation	Supervisor line of sight; device lockers at entry	Written policy plus agent attestation
Clean desk in regulated environments	Visual inspection; no paper permitted on the floor	Written policy plus agent attestation
No unauthorized person able to see or hear	Controlled physical access to the operations area	Written policy plus agent attestation

The pattern is identical in each case. A control that was *observed* became a control that is *promised*. This is arguably the largest uncontrolled change in contact center risk posture of the past six years, and most risk registers have not been updated to reflect it.

Converting promises back into observations

Platforms in this category can detect and alert on a mobile device appearing in frame, a second face appearing in frame, and an unattended workstation, in real time, to a supervisor. Where a wide angle camera is used, clean desk conditions can be assessed as well.

The effect is to turn three attestations back into three observations, which is the only form in which they were ever actually controls. Note that this is the same mechanism as everything in Section 5, pointed at a different problem. It is not a separate product and it should not be a separate business case.

What to require during diligence

- Current PCI DSS and SOC 2 Type II attestation reports, provided rather than asserted. Check the dates and the scope.
- A written statement of what is retained. The correct answer for this category is nothing.
- Confirmation of which countries your program operates in, including any your BPO subcontracts to. Subcontracting disclosure deserves a direct question of its own.

> **A regulatory note**
>
> Requirements differ materially by jurisdiction, and remote work provisions in regulated industries have moved quickly. Several US states have introduced legislation governing remote handling of consumer financial data. Verify current obligations for every country your program touches rather than relying on a vendor summary, including this one.

> **A control that was observed became a control that is promised. Nothing else about your risk posture changed as much, and almost nothing gets less attention.**

SECTION 7

What the evidence shows

Measured results from four contact center deployments of a virtual floor platform currently in the Cloud Tech Gurus portfolio. Customer names are withheld. Industries and measurement windows are stated. Unless noted, figures were measured within thirty days of production go live, against the operation's own prior baseline.

Operation	Window	Measured change
Receivables management firm, US	30 days	Daily production hours per agent **6.5 to 7.4**. Calls taken **+33%**. Revenue **+27%**. Positive return inside 30 days.
Receivables management firm, US	30 days	Call volume **+51%**. Call handling **+26%**. Cash collections **+8.3%**.
Accounts receivable services, US	30 days	Call volume **+12.66%**. Revenue per hour **+8.04%**. Daily productivity **+7.19%**.
Vehicle protection products provider	4 months, vs. benchmark team	Abandonment rate **−32%**. After call work **−24%**. Monthly call volume **+10.5%**. Average hold time **−3.7%**.

In every case headcount did not change. No new dialer was purchased, no platform was replaced, and no additional supervisors were hired. The variable was whether supervisors could see and help their people.

The production hours figure in the first row is the one we would put in front of a CFO. Moving 6.5 hours to 7.4 at constant payroll is a fourteen percent gain in paid hour yield, and your finance team can test it against their own agent count in about five minutes.

Worth noting what these numbers are not. None of them came from catching anyone doing anything wrong. They came from agents getting answers faster, supervisors intervening earlier, and new hires reaching competence sooner.

Section 7: What the evidence shows

The limits of this evidence

We would rather you hear this from us than discover it in a diligence call.

- **All four operations were managing their own remote or blended agents, not a third party's.** The outsourced application described in this briefing is a newer use of the same capability, and there is not yet a named, published case study for it. What exists is the platform, the results opposite, and a growing number of deployments where a brand has extended its own floor to include a partner's agents.
- **Three of the four operate in receivables.** That sector has unusually direct revenue attribution, which makes gains easy to measure and may make them look larger than they would in a service only environment. Treat the vehicle protection figures, meaning abandonment, after call work and hold time, as the better analogue for a customer service operation.
- **These are supplier provided case studies** measured by the operating companies themselves, not independently audited. We have reviewed the methodology and consider the figures credible. We have not verified them line by line, and we will say exactly that to your CFO.
- **Thirty day windows favour the intervention.** Novelty effects are real. Ask any vendor in this category for twelve month retention and performance data, and treat reluctance as an answer.

> **The strongest argument for this category is not in the table opposite. It is on the next page, and it was never a product requirement.**

Six people

Across roughly one hundred deployments in eighteen countries, six agents working remotely have received emergency medical or police response because a supervisor watching the virtual floor saw what was happening and called for help. Two cardiac events. Two strokes. An anaphylactic reaction. A home invasion.

None of that was a design goal. There is no feature named for it, no line in a specification, and the supplier is careful to say it is not a reason to buy software. We agree. It is in this briefing for two other reasons.

The first reason is arithmetic

Those six reports came from two customers out of approximately one hundred, and four of the six came from a single customer running around three hundred licenses. The likeliest reading is not that six is a remarkable number. It is that six is the number that happened to get reported back to a vendor, and the real figure across a hundred deployments is higher.

The second reason is what it says about distance

Every one of those six people was having a medical or personal emergency while at work, alone, in a job that used to be done in a room full of colleagues. On a floor, someone would have seen within seconds. Distributed, with nothing in place, nobody sees at all.

That is the real subject of this document, and it is easy to lose inside a discussion of abandonment rates and quality sampling. The reason to be able to see your agents is not to verify that they are working. It is that people who are alone at work are genuinely alone, and some of the consequences of that are not measured in average handle time.

> **Where this applies to your outsourced seats**
>
> Ask yourself who would notice if this happened to an agent on your program in Manila or Johannesburg tomorrow. Then ask whether your answer would be acceptable if the same agent sat in your building. Those agents are working on your behalf, and the standard of care you are able to extend to them is a decision you are making whether or not you have made it deliberately.

SECTION 8

How to evaluate this category

Ten questions we put to every supplier in this category, reproduced without modification so you can use them independently, including against platforms we do not represent and against the one we do.

1. **Can a supervisor see more than one agent screen at a time?** Most conferencing and collaboration tools permit a single screen share. If the platform cannot mirror multiple monitors across multiple agents simultaneously, it is not replicating a floor. This is the single most clarifying question in the list.
2. **Is anything recorded or retained?** Ask specifically about video, screen capture and personal data, and get the answer in writing. A platform that records creates a discovery liability and a works council problem the moment you deploy it.
3. **Can the agent initiate?** Can they raise a hand, message a supervisor, call a peer, see who is present? If help only flows one direction you have bought monitoring software and your agents will treat it accordingly.
4. **Can a supervisor hear the live customer call while watching the screen?** Observation without call audio produces guesses about what an agent is doing wrong.
5. **What does new hire nesting look like on the platform?** Ask for a walkthrough of a trainee shadowing a tenured agent. This is where the strongest operators find unexpected value, and where weaker platforms have nothing to show.

Section 8: How to evaluate this category

Questions six through ten

6. **Which certifications are current, and can you see the reports?** PCI DSS and SOC 2 Type II attestations should be provided, not asserted. Check the dates and the scope.
7. **Does it work through a browser?** A desktop agent creates a deployment dependency on your BPO's IT organization, which is frequently the thing that kills the project.
8. **Will your BPO consent, and who pays?** This is the real commercial question. In our experience BPOs value these platforms operationally but resist capitalizing them, preferring the client to own the licenses. Decide your position before you open the conversation, and write it into the next contract renewal rather than negotiating mid term.
9. **What is the contract term, and how fast is a pilot?** Month to month agreements exist in this category, and two weeks to a live pilot is achievable. A vendor requiring a multi year commitment for a tool this straightforward to deploy is pricing in your switching cost, not their delivery risk.
10. **What is the documented evidence for your specific use case?** Ask for the case study that matches your application, not the vendor's strongest one. If it does not exist, a credible vendor will tell you so. Ours did, and you can read the result in Section 7.

> **On question ten**
>
> Question ten is the one most likely to expose a weak recommendation, including ours. It is in the list deliberately. An evaluation process that cannot survive its own rubric is a sales process wearing a rubric's clothing.

Take this list into any vendor conversation. If a supplier objects to the questions rather than answering them, that is useful information, and it did not cost you anything to obtain.

SECTION 9

What we recommend, and why we are saying it

The capability described in this briefing exists today. It is patented, deploys through a browser in one to two weeks, is available on month to month terms, holds current PCI DSS Level 2 and SOC 2 Type II attestations, and records nothing. It is in use across approximately one hundred contact centers in eighteen countries. It sits in the Cloud Tech Gurus portfolio and we will introduce you to it directly.

We have not named it here, for a reason worth being open about. Lead with a product name and the conversation becomes a demo, and you lose the part with actual value to you: a neutral read on whether this category solves your problem at all, and which platform in it fits your environment. Some operators who read this will conclude the answer is no. That is a legitimate outcome and we would rather reach it in twenty minutes than in a procurement cycle.

What Cloud Tech Gurus is

A CX consulting and advisory firm focused exclusively on contact center and customer experience. We maintain relationships with more than 250 suppliers across 58 solution categories and more than 40 BPO organizations, which is what lets us shortlist against a category rather than sell a product. Our advisors are practitioners: former contact center directors, VPs and SVPs who operated these environments before they evaluated them. We did not study contact centers. We ran them.

250+
Suppliers under evaluation across the CX technology landscape

58
Solution categories, each with a documented evaluation rubric

40+
BPO relationships informing the outsourcing perspective here

How we are paid, in plain terms

Suppliers pay us. You do not. Sourcing, shortlisting and evaluation support are provided at no cost to the buying organization under a Letter of Agency. That model carries an obvious conflict, so we manage it the only way that holds up: a documented, auditable rubric per category, applied consistently, with the evidence and its limits disclosed to you, as in Section 7.

Section 9: What we recommend

Three ways to take this further

- **Score your own program.** Use the four tier model in Section 3. If you land in Tier 1 or Tier 2 and hold regulated data on distributed screens, that combination deserves a conversation this quarter rather than next year.
- **Ask for the unscheduled hour.** One hour, six agents, no notice. Frame it as wanting to see how your partner's people are supported, because that is what you are actually asking. Whatever the answer, you will know more than you did this morning and it will cost you nothing.
- **Ask us for the four case studies.** We will send the full breakdowns including methodology notes and caveats, and tell you which of the four resembles your operation and which does not.

If none of those appeals, the four tier model and the ten questions are yours to keep and use with anyone. That is why we published them.

> **Your agents represent your brand to your customers. All of them, including the ones on someone else's payroll.**

Talk to us

To discuss any of this, or to ask for the four case studies, reach out to **Darren Prine**, Partner and Chief Revenue Officer, at darren@cloudtechgurus.com, or contact the team at info@cloudtechgurus.com.

Cloud Tech Gurus · Phoenix, Arizona · cloudtechgurus.com

No cost, no obligation, and no vendor introduction unless you ask for one.

Performance figures are drawn from supplier provided case studies measured by the operating companies named therein, reviewed but not independently audited by Cloud Tech Gurus. Certification, deployment and contract term details reflect supplier representations current as of August 2026 and should be verified during diligence. Regulatory requirements vary by jurisdiction. This briefing is informational and is not legal, regulatory or financial advice.